

**RADIO COMMUNICATION SYSTEM EMPLOYING  
SPECTRAL REUSE TRANSCEIVERS**

**CROSS-REFERENCE TO RELATED APPLICATION**

[001] The present invention claims the benefit of co-pending U.S. Application Serial No. 60/432,223 filed December 10, 2002, by Edward Gerhardt et al, entitled: "Link Utilization Mechanism for Secondary Use of a Radio Band," assigned to the assignee of the present application and the disclosure of which is incorporated herein.

**FIELD OF THE INVENTION**

[002] The present invention relates in general to communication systems, and is particularly directed to a link utilization control mechanism for secondary users in a licensed band, or users in an unlicensed band, which employs a spectral reuse transceiver (SRT) operating with a selectively filtered form of orthogonal frequency division multiplexing (OFDM) for producing non-interfering radio channels, particularly, VHF and UHF radio channelized communications bands.

**BACKGROUND OF THE INVENTION**

**[003]** In some radio bands, for example, the 217-220 MHz VHF band, the U.S. Federal Communications Commission (FCC) grants non-exclusive primary licenses use of the band for various communication services, such as push-to-talk voice. Primary users (or PUs) pay for these non-exclusive rights with the expectation that they will not suffer interference by other users. The FCC also allows secondary users (or SUs) to use the same band and the same channel within the band on a 'non-interfering' or secondary basis. Namely, so long as the primary user is not using a channel within the band, a secondary user may use it. Similar policies are also practiced by frequency management authorities in foreign countries.

**[004]** The FCC and similar agencies in foreign countries are anxious to find ways to enable expanded use of such radio frequency bands, without reducing the quality of service available to the licensed primary users. For secondary users, this band - and other similar bands - represent a cost-free band with excellent radio transmission properties for telemetry and other uses. Because secondary users must not interfere with primary users, complaints of interference from a primary user to the FCC may result in an administrative order from the FCC requiring that the secondary user move to another portion of the band or leave the band entirely.

**[005]** This movement can be disruptive to the secondary user's service and can be expensive, especially if site

visits, equipment modification, or exchange is required to implement the change. Therefore, a mechanism is needed by which a secondary-use radio may employ the band on a non-interfering basis and which will adapt the radio's frequency usage should new primary users appear. It should be noted that primary users (who are licensed) always have priority over secondary licensed users or secondary unlicensed users; there is no first-use channel frequency right for secondary users.

**[006]** Interference within such a band can occur in at least two ways. First, the primary user and one or more secondary users may have occasion to use the same channel and significantly disrupt one another. Secondly, one or more secondary users may transmit briefly with enough power and for a sufficient duration to cause some unwanted effect on the primary user's radio, such as a burst of noise through it's speaker or loss of data. Both of these and other types of interference must be avoided.

**[007]** Primary users typically are field personnel who use an assigned channel for occasional voice transmissions to coordinate their activities with a central site or with co-workers in a nearby locale. It is likely that band usage will increase in the future. The aforementioned disruption of service caused by the appearance of new primary users warrants the need for a communication technique that will effectively avoid interference by secondary users.

**[008]** Because of the occasional use of the band by primary users, it has been difficult to directly observe what frequencies are being used. It has also been difficult to coordinate frequency use by monitoring applications for primary use because of administrative delays. A further difficulty is that manual coordination must be done continuously, since new primary users may appear at any time.

**[009]** One currently employed approach for having secondary users share a frequency band is to use manual frequency coordination. In accordance with this scheme, an installer or maintainer of a secondary user of a narrowband or wideband radio selects a frequency that appears not to be used, and then manually adjusts the radio through physical controls, or through a management control channel (in-band or out-of-band). The frequency selection process is typically performed by examining primary user license applications, or through the use of an instrumented spectral survey. The installer or maintainer then selects a frequency with sufficient, contiguous unused bandwidth to support the application. One problem with this approach is that it is labor-intensive, and it is difficult to differentiate between primary users and other secondary users.

**[010]** An additional problem with this frequency selection process is that a primary user may appear only occasionally or a new primary user may appear unexpectedly. Without a further means of frequency coordination, the secondary user may interfere with the

primary user and become subject to an administrative order to change the frequency or exit the band, as noted above. Another problem with this approach is that, in a busy band, it may become difficult to find enough contiguous bandwidth to support a given application. Moreover, adapting manual frequency radios to new services that require different or varying bandwidths may be difficult or expensive, depending on how bandwidth of the radio is adjusted.

**[011]** Another technique for frequency coordination is to monitor applications for primary use licenses. This solves the difficulty of differentiating primary users from secondary users, but suffers all the other problems associated with human monitoring of spectrum through instrumentation, described above.

**[012]** A further approach for allowing secondary users to share a frequency band is through the use of agile frequency, narrowband or wideband radios with automatic frequency coordination. This type of radio automatically detects interference from other radios and adjusts its frequency. For example, if the radio detects a high bit error rate for a sustained period, it may begin scanning for other, available frequencies or use prescribed alternate frequencies. While this constitutes an improvement over manual frequency coordination processes, it suffers from some of the same issues associated with manual frequency coordination. For example, it may be difficult to find contiguous bandwidth for the application, it may be difficult to

distinguish between primary and secondary users, or it may be difficult to adapt the radio to new services that require different or varying bandwidths. Another problem with this approach is that it interferes with primary users during the interference detection phase and possibly during the frequency change phase, depending upon implementation.

**[013]** Still another approach for enabling secondary users to share a frequency band with primary users is through the use of a frequency hopping spread spectrum (FHSS) radio. This type of radio spreads the transmission over a very wide band via frequency hopping with low spectral density. The radio occupies the entire band, but operates at a level sufficiently close to the noise floor to minimize the effect that the transmission will have on narrowband or wideband radios. Such a radio intends to solve the frequency coordination problem of manual frequency or agile frequency radios described above, and has the added benefit of coding gain.

**[014]** However, a problem with this type of radio is the fact that the spread spectrum modulation can produce unknown effects on narrowband and wideband radios, such as possible squelch circuit activation or other types noise caused by interference. Another problem is that above 1 GHz (such as in the 2.4 GHz band), it may be difficult to close the radio link at the distances needed for many applications, such as telemetry using cost-effective radio equipment. Another problem is that when two secondary (licensed/unlicensed) users are

communicating, there may be a higher-powered, interfering primary user that is closer to the secondary user transmitter than the intended secondary user receiver. In this case, the power in the FHSS signal must be increased to overcome the strong interference, which increases interference to the primary user, and decreases the jamming margin. (This latter problem is often referred to as the 'near/far' problem.) An additional problem is the fact that FHSS inherently occupies all channels, including primary user channels; therefore, it is not clear how the Federal Communications Commission in the United States, and equivalent authorities in foreign countries, will view FHSS for secondary use.

#### **SUMMARY OF THE INVENTION**

[015] In accordance with the present invention, the problem of how to allocate secondary usage of a licensed radio frequency band (e.g., the 217-220 MHz band comprised of 480 - 6.25 KHz sub-bands or channels), in a manner that will not interfere with other (primary) users is successfully addressed by a new and improved spectral activity-based link utilization control mechanism and efficient modem algorithm. Pursuant to this mechanism, each site contains a spectral reuse transceiver, that operates with a selectively filtered form of orthogonal frequency division multiplexing for producing a sub-set of non-interfering radio channels,

particularly, VHF and UHF radio channelized communications bands.

**[016]** The spectral reuse-based communication network of the invention is that of a star-configured system, wherein a spectral reuse transceiver installed at a master or hub site communicates with spectral reuse transceivers installed at multiple remote sites. The master site's spectral reuse transceiver may employ an omnidirectional antenna, so that the master site may communicate with any of the remote site spectral reuse transceivers, while each remote site spectral reuse transceiver may employ a directional antenna whose boresight axis is pointed toward the master site.

**[017]** Pursuant to the invention, the master site spectral reuse transceiver periodically initiates a clear channel assessment routine, in which the master site and each of the remote sites participate. This routine is used to compile an aggregate list of non-interfering or 'clear' channels (e.g., a plurality of 6.25 KHz wide channels) that may be employed by participants of the network. By transmitting on only what has been previously identified as a clear channel, a respective site's spectral reuse transceiver is ensured that it will not interfere with any primary user of the (3 MHz wide) band of interest.

**[018]** Except when it is transmitting a message to the master site, each remote user site sequentially steps through and monitors a current list of clear channels (that it has previously obtained from the master site), in accordance with a pseudo random hopping sequence that



is known a priori by all the users of the network, looking for a message that may be transmitted to it by the master site transceiver.

**[019]** When a remote site transceiver is not transmitting, it is listening for ticklers from the master SRT. This does not require a sequential stepping through the clear channel list. The position of the tickler carriers is known a priori and are derived from the clear channel list. The only time that the remote site's transceiver sequentially steps through the channels is either during clear channel assessment or during the join (network discovery) process.

**[020]** During the preamble period of any message transmitted by the master site, each remote site's transceiver scans all 480-6.25 KHz frequency bins within the 217-220 MHz spectrum for the presence of energy. Any bin containing energy above a prescribed threshold is marked as a non-clear channel, while the remaining channels of the 480 possible are identified as clear channels.

**[021]** The master site transceiver then sequentially interrogates each remote site transceiver for its clear channel list via a clear channel request message. In response to receiving a clear channel request message, a respective remote site transceiver transmits back to the master site transceiver the clear channel list it obtained as it scanned the 3 MHz band during the preamble portion of the master site's message. Once the master site has completed its interrogation of all the

remote sites, it logically combines all of the clear channel lists it has received from all the interrogated remote sites, to thereby produce an 'aggregate' clear channel list. This aggregate clear channel list is stored in the master site's transceiver and is broadcast to all of the remote sites over a subset of the clear channels that is selected in accordance with a PN sequence through which clear channels are selectively used among the users of the network. The aggregate clear channel list is normally transmitted to the remotes using the normal multi-carrier frequency hopping transmission. Only the initialization message will transmit the clear channel list in a single-carrier fashion. When the aggregate clear channel list is received at a respective remote site it is stored in its transceiver.

**[022]** A synchronization function is carried out in both the master and remote transceivers and ensures that the endpoint radios are synchronized to a common set of the clear channel list, channel hopping sequence, and preamble channel. A remote transceiver determines a lack of network synchronization by an absence of received network messages. Primarily, a remote transceiver will use the expected periodic aggregate clear channel list as a network "heartbeat." The remote transceiver will also maintain a counter indicating the amount of time elapsed from the last network message received. Then based upon the activity level of these two indicators, the remote transceiver will determine whether it is

properly synchronized with the network. If the remote transceiver determines that it is not properly synchronized, it will fall into the normal initialization sequence for an uninitialized node. In the special case that a remote transceiver detects a primary user on the preamble channel, the remote transceiver will send a message to indicate the conflict such that the master transceiver can indicate to the remaining remote transceivers to fall back to the next preamble channel candidate based on a predetermined algorithm. The remote transceiver will then switch to the next preamble candidate as well. In the case where the master transceiver detects a primary user on the preamble channel, it will initiate the switch to the next preamble channel candidate by sending a message indicator to all remote transceivers. Those transceivers not in conflict with the primary user will receive the message and switch. Those remote transceivers that are in conflict will likely be unable to receive the message and will ultimately time out and declare a loss of network synchronization. At this point they will initiate the normal initialization sequence.

**[023]** In order to enable the master and remote sites to successfully communicate with one another over 'clear' secondary use channels in the manner described above, the master and remote users employ a spectral reuse transceiver architecture, that operates using a selectively filtered form of orthogonal frequency division multiplexing for producing a sub-set of

non-interfering (6.25 KHz) channels. The spectral reuse transceiver of the invention includes a transmit signal processing path and a receive signal processing path.

**[024]** In the transmit path through the spectral reuse transceiver, a respective digital data packet from a data communication source is interfaced with a front end packet buffer. As digital data is read out from the packet buffer it is initially subjected to forward error correction, as by way of a Viterbi encoder, and then coupled to a modulator. Pursuant to a non-limiting, but preferred embodiment of the invention, the data modulation scheme is differential quaternary phase shift keyed (DQPSK) modulation. DQPSK has been chosen as a preferred modulation, since it avoids the need for carrier phase lock and non-coherent demodulation can be used. Also, in an OFDM system, phase uncertainty can be caused by either carrier phase offset or sample timing offset. A sample timing offset produces a phase offset on the carriers that is linearly increasing with the carrier number. DQPSK relieves some of the sample timing synchronization requirements. In addition, synchronization in most OFDM systems utilize pilot carriers during acquisition and tracking. It is preferable not to use such techniques in the spectral reuse transceiver of the present invention, which frequency hops only a limited number of data bearing carriers. This makes synchronization more complex. DQPSK removes these problems.

**[025]** In-phase (I) and quadrature phase (Q) outputs of the DQPSK modulator are coupled to a symbol mapper. The symbol mapper maps the I and Q data from the modulator into appropriate 'clear' channel carrier(s) using a prescribed hopping PN sequence known a priori to all the users of the network. The clear channel map is a compilation of all currently 'not being used', or 'clear' channels within the band of interest, namely which of the 480-6.25 KHz spaced channels that make up the band from 217 MHz to 220 MHz have been determined during the clear channel assessment routine, described above, to be effectively 'unused' and thus available for secondary use. The spectral reuse transceiver's communication control processor stores the clear channel map that is used by the symbol mapper. The clear channel hopping sequence is determined by means of a pseudo-random sequence defined by a generator polynomial and a starting PN seed value. For each pair of I/Q data retrieved from the DQPSK modulator, a symbol map look-up table is read out, to locate the appropriate carrier for the data. The value serves as an index into an array that is used to hold a multi-sample input vector for an inverse Fourier transform unit of a modulated filter bank.

**[026]** Prior to being coupled to the modulated filter bank, the mapped carrier outputs of the symbol mapper are coupled to a gain multiplier, which operates on a per transmission burst basis, and serves to minimize the power level necessary to complete the radio link between

the master and a remote site, and thereby avoids interfering with other links. The modulated filter bank is comprised of an inverse FFT unit and an associated polyphase filter. In accordance with the parameters of the exemplary embodiment described herein, the present invention operates within the 217-220 MHz frequency band, with adjacent carrier spacings of 6.25 KHz, yielding a total of 480 available carriers. However, among this substantial number of available frequencies only a relatively small fraction of the carriers (e.g., on the order of a dozen to several tens of channels) may actually be used.

**[027]** In order to accommodate the 480 carriers within the 217-220 MHz band of interest, the inverse FFT unit is preferably configured as a 512 point inverse FFT. The polyphase filter is designed from a low-pass filter prototype, formed as a bank of bandpass filters, which are frequency-shifted versions of the low-pass filter prototype. In accordance with the invention, the low-pass filter prototype function is implemented as a six symbol wide interpolated or over-sampled, root-raised cosine filter, that is multiplied by a Tukey window (with  $\alpha = 5$ ) to further reduce sidelobes. As will be described, the low-pass impulse response of such a windowed filter mechanism settles rapidly to a relatively flat characteristic on either side of the main lobe, making it particularly suitable for use with an inverse FFT. Such a windowed filter not only has

superior sidelobe performance, but retains near zero inter-symbol interference (ISI).

**[028]** A respective filter section of the interpolating polyphase filter in accordance with the present invention comprises a six stage delay line, respective  $z^{-1}$  stage outputs of which are coupled to coefficient multipliers, whose outputs are summed to realize an interpolated carrier value. In the exemplary parameters of the invention, the interpolation factor is times 1.5, so that for a  $M=512$  point inverse FFT, a commutating switch coupled to the respective filter stages steps through or is commutated among the 512 points of the FFT 768 times (i.e., for each symbol it steps through the 512 inverse FFT output samples 1.5 times), producing 768 samples.

**[029]** Thus, with the 512 complex values produced by the IFFT placed into six individual element delay line ( $z^{-1}$ ) stages, the contents of a given delay line form a vector of six complex elements, and there are 768 polyphase filter components computed from the prototype filter. When the commutator selects a given inverse FFT sample, the contents of a respective delay line are multiplied by the component filter (i.e., as a vector dot product) and then summed together to form one output sample of the polyphase filter.

**[030]** The output of the polyphase filter is interpolated and then translated up to an intermediate frequency of 70MHz. The data stream is then filtered and applied to digital-to-analog converter (DAC). The output of DAC is

translated via a frequency translator to the desired 217-220 MHz transmission frequency band for transmission.

**[031]** The receive signal processing path direction through the spectral reuse transceiver is essentially complementary to the transmit path, described above. In particular, a received signal in the 217-220 MHz transmission band is coupled to frequency translator, which down-converts the received signal to IF. The output of the frequency translator is coupled to an analog-to-digital converter (ADC), which produces a baseband digitized output that is applied to a digital mixer. The digital mixer down-converts the received signal further, in accordance with the output of a synthesizer and applies the down-converted output to a decimator unit, whose output is the same frequency as output by the polyphase filter of the transmit section of the transceiver.

**[032]** The baseband data stream is buffered in a timing correction buffer and read out under the control of a symbol timing estimator and applied to a polyphase filter, the output of which is coupled to an FFT unit. The output of the FFT unit is coupled to a frequency offset estimator, whose output is coupled to coarse and fine frequency estimation synthesizers. Coarse frequency estimation serves to reduce frequency offset to within a prescribed fraction (e.g., one-half) of adjacent carrier spacing; fine frequency estimation is performed on a burst-by-burst basis on the pure carrier portion of the



preamble only. Fine frequency estimation is performed by looking at the change in phase of the carrier from symbol to symbol. The change in phase over four symbols is used as the fine frequency offset estimate.

**[033]** The output of the FFT unit is also coupled to a symbol de-mapper. The symbol de-mapper employs the same mapping table as the symbol mapper in the transmit path section, described above, so that the frequency data may be recovered. The FFT is also coupled to a received signal strength indicator, which feeds the symbol timing estimator. The symbol timing estimator is also coupled to the output of a mixer which multiplies the received signal by the output of a preamble synthesizer which downconverts the preamble channel to baseband. The symbol timing estimator processes this baseband data.

**[034]** The first part of the preamble of the received signal contains only pure unmodulated carrier, so as to provide an indication that the preamble is beginning. Detecting energy in that frequency bin of the FFT of the receiver means that the receiver needs to be monitoring a burst, and the output of the received signal strength indicator (RSSI) is used to enable the symbol timing estimator to find the boundaries of a received symbol. Once the symbol boundaries have been determined, the symbol timing estimator causes an appropriate shift of the data in a data buffer, to apply the received data one symbol at a time into a polyphase filter. The de-mapped frequency representative I and Q data values from the symbol de-mapper are coupled to a DQPSK demodulator.

The output of the DQPSK demodulator is coupled to a Viterbi decoder, which recovers the original data and couples the data to the packet buffer.

**[035]** The spectral reuse transceiver of present invention also transmits a plurality of 'tickler' tones (e.g., three to five a priori known carriers) within the (217-220 MHz) band of interest simultaneously over a short duration (e.g., four symbols, wherein a symbol is a group of 768 complex samples output by the transmitter). These tickler tones are used to initiate prescribed actions in the receiving spectral reuse transceiver. Each set of tickler tones is different from each other set and defines the action to be taken by the other transceiver.

**[036]** Since the channels which constitute the tickler carriers are known, the frequency information for those (three) channels is extracted from the FFT. The absolute values of the I and Q sub-channels are summed and then normalized using an AGC value as derived from an AGC loop. The normalized value is subjected to a moving average over a prescribed symbol width (e.g., four symbols). The resulting average is then compared with a threshold. If the threshold is exceeded an output is triggered indicating the detection of the tickler tones.

**[037]** As noted above, during the preamble phase of a transmitted burst, only the preamble carrier contains transmitted power from a spectral reuse transceiver. During this time, a clear channel assessment (CCA) operation is conducted, wherein each transceiver

monitors the 217-220 MHz band for the presence of energy. (such as that sourced from primary users). During the preamble, after RSSI and frequency offset computations, most of the remaining preamble is used for symbol timing recovery. During this time the primary goal of each transceiver's control processor is to monitor the 480-6.25 KHz channels in the 217-220 MHz band for the presence of interferers (primary users and others).

**[038]** Clear channel assessment is performed at both the master site and at each of the remote sites in order to generate a current list of non-interfering (clear) channels that is compiled by the master site and distributed thereby to all remote sites in the network. By transmitting on only a clear channel, a transceiver is ensured that it will not interfere with any primary user of the spectrum of interest.

**[039]** As noted previously, each remote site transceiver conducts communications with the master node and the master node conducts communications with any of the remote nodes. For this purpose the master site may employ an omnidirectional antenna, while the remote sites may use directional antennas aimed at the master site. Since, for the most part, remote sites cannot effectively listen to each other's transmissions, a virtual carrier sensing routine is carried out to convey a pending access to the channel. This mechanism is employed to negotiate multiple access of the channel. Virtual carrier sensing is defined by a prescribed

pattern of clear channels that are active for a prescribed number of symbol periods (e.g., five).

**[040]** The remote sites look for this pattern of clear channels and do not attempt to access the channel. A different patterned carrier sense transmission is used to indicate that the channel is clear for access requests. Each remote site transceiver with a pending message awaiting transmission will respond through a random slotted back-off, before attempting a channel access request. Once a remote site has been granted access to the channel, the master site listens for a potential transmission from the remote site during an acquisition of signal (AOS) timeout period. If a prescribed time-out occurs before a request access from a remote site transceiver is received, it is inferred that there are no remote transceivers awaiting transmission of a message. The master will only send a master access if it has a pending message to send. On the other hand, if a remote site transceiver has a message awaiting transmission, the time-out will not expire before a request transmission is received from a remote site transceiver. The master site transceiver then grants access to the remote site, followed by receipt of the remote site's message. Once the remote site's message has been received, the master site transceiver proceeds to transmit ticklers and beacon preambles at prescribed intervals or messages as they become available.

[041] From a remote site transceiver's standpoint, the remote site transceiver is initially listening for a media open tickler or a master access tickler. When a media open tickler is received, the remote site transceiver delays transmitting, or 'backs off' a random number of time slots, and then transmits a request access tickler (unless, during the back-off period, the remote receives an access grant tickler from the master, which indicates that some other remote has gained control of the media). When the remote site transceiver receives an access grant tickler from the master site, it transmits a message before returning to its monitoring state. Where a master access tickler is received, the remote site transceiver listens for a message from the master. If the message from the master is not received within a timeout period, then the remote will return to listening for ticklers.

#### **BRIEF DESCRIPTION OF THE DRAWINGS**

[042] Figure 1 diagrammatically illustrates the overall architecture of a communication network, respective terminal unit transceiver sites of which employ the spectral reuse transceiver of the invention;

[043] Figure 2 diagrammatically illustrates the architecture of a spectral reuse transceiver in accordance with the present invention;

[044] Figure 3 is a dual graphical plot of the theoretical bit error rate (BER) curve for DQPSK compared with the variation for coherent QPSK;

**[045]** Figure 4 shows a constellation for  $\pi/4$  DQPSK;

**[046]** Figure 5 depicts the multicarrier frequency spectrum produced by an inverse fast Fourier transform absent a polyphase filter;

**[047]** Figure 6 diagrammatically illustrates a mathematical model of a communication system employing a modulated filter bank structure, which can be readily implemented with a discrete Fourier transform (DFT)-based signal processing operator (e.g., an IFFT) followed by a polyphase filter;

**[048]** Figure 7 shows the windowed lowpass prototype filter impulse response (the basis of the polyphase filter);

**[049]** Figure 8 shows the composite spectral response of a modulated filter bank for a pair of adjacent carrier frequencies;

**[050]** Figure 9 diagrammatically illustrates the architecture of an interpolating or over-sampled polyphase filter in accordance with the present invention;

**[051]** Figure 10 plots bit error rate vs.  $E_b/N_0$  for a demodulator employed in the spectral reuse transceiver of the invention where there is a sample timing error;

**[052]** Figure 11 shows the composition of a burst preamble employed for the spectral reuse transceiver of the invention;

**[053]** Figure 12 shows the signal processing flow path through a symbol timing estimator for a spectral reuse transceiver of the invention;

[054] Figure 13 shows the output of the moving average correlator of the symbol timing estimator of Figure 12;

[055] Figure 14 plots symbol timing estimate in the absence of adjacent channel interference;

[056] Figure 15 plots symbol timing estimate in the presence of adjacent channel interference;

[057] Figure 16 shows the spectral transfer function of a forty tap implementation of the lowpass filter used for adjacent channel interference rejection in the symbol timing estimator of Figure 12;

[058] Figure 17 is a block diagram of a received signal strength indicator algorithm;

[059] Figure 18 graphically plots the output of the RSSI algorithm of Figure 17;

[060] Figure 19 graphically plots the output of the RSSI algorithm when AGC normalization is not performed;

[061] Figure 20 diagrammatically illustrates the configuration of a front end AGC unit;

[062] Figure 21 depicts a time domain waveform of a received signal in the presence of adjacent channel interference;

[063] Figure 22 shows the effect of AGC attenuation on an adjacent channel interferer;

[064] Figure 23 is a graphical plot of bit error rate vs.  $E_b/N_0$  in the presence of frequency offset at the receiver;

[065] Figure 24 diagrammatically illustrates the signal processing architecture of a fine frequency estimator;

[066] Figure 25 shows the accuracy of the fine frequency estimator for an  $E_b/N_o$  of 5 dB when averaging over four pure carrier symbols;

[067] Figure 26 shows the signal processing flow path for tickler detection;

[068] Figure 27 is a state diagram of the general operation of the spectral reuse transceiver of the present invention;

[069] Figure 28 shows the contents of a 'beacon preamble' burst, that is periodically transmitted by a master site transceiver;

[070] Figure 29 shows the contents of an 'initialization' single-carrier burst which is transmitted by a master site transceiver;

[071] Figure 30 shows the contents of a standard message burst used by a master site transceiver for the transmission of information to a remote site transceiver;

[072] Figure 31 shows a 'tickler' burst;

[073] Figure 32 is a state diagram illustrating the manner in which a remote site transceiver 'joins' (or reacquires) the network;

[074] Figure 33 is a sequence diagram of a clear channel assessment routine through which a master node maintains and distributes an aggregate list of clear channels;

[075] Figure 34 is a state diagram associated with the operation of a remote site transceiver;

[076] Figure 35 is a sequence diagram associated with the state diagram of Figure 34;



[077] Figure 36 is a contention and backoff diagram associated with the sequence diagram of Figure 35; and

[078] Figure 37 is a state diagram associated with the operation of a master site transceiver.

#### **DETAILED DESCRIPTION**

[079] Before detailing the spectral reuse transceiver for an OFDM-based communication system and associated link utilization control mechanism of the present invention, it should be observed that the invention resides primarily in a prescribed arrangement of conventional communication circuits and components, and supervisory digital control and signal processing circuitry that controls the operations of such circuits and components. Consequently, in the drawings, the configuration of such circuits and components, and the manner in which they are interfaced with various communication circuits have, for the most part, been illustrated by readily understandable block diagrams, which show only those specific details that are pertinent to the present invention, so as not to obscure the disclosure with details which will be readily apparent to those skilled in the art having the benefit of the description herein. Thus, the block diagrams of the Figures are primarily intended to show the various components of the invention in convenient functional groupings, so that the present invention may be more readily understood.

[080] Attention is initially directed to Figure 1, which diagrammatically illustrates a master site - multiple

remote site configured communication network, respective terminal unit transceiver sites of which employ the spectral reuse transceiver of the invention, to be detailed below. As shown therein the network contains a master terminal transceiver site or hub 10 and a plurality of remote terminal transceiver sites 12. As pointed out above, the master site's spectral reuse transceiver 10 may employ an omnidirectional antenna, so that the master site's terminal unit may communicate with any of the remote site transceivers, while each remote site spectral reuse transceiver 12 may employ a directional antenna whose boresight axis is pointed toward the master site. The remote sites are intended to communicate exclusively with the master site.

**[081]** As will be detailed below, the master site periodically initiates a clear channel assessment routine, which is conducted at both the master site 10 and at each of the remote nodes 12. This routine is used to gather a list of non-interfering or 'clear' communication channels that may be employed by participants of the network. By transmitting on only what has been identified as a clear channel, a respective site's transceiver is ensured that it will not interfere with any primary user of the 217-220 MHz band of interest.

**[082]** When a remote site transceiver is not transmitting, it is listening for ticklers from the master SRT. This does not require a sequential stepping through the clear channel list. The position of the tickler carriers is

known a priori and are derived from the clear channel list. The only time that the remote sequentially steps through the channels is either during clear channel assessment or during the join (network discovery) process.

**[083]** During the preamble period of any message transmitted by the master, each remote site's transceiver scans all 480-6.25 KHz frequency bins within the 217-220 MHz spectrum for the presence of energy. Any bin containing energy above a prescribed threshold is masked as a non-clear channel, while the remaining channels of the 480 possible are marked as clear channels.

**[084]** The master site transceiver 10 sequentially interrogates each remote site transceiver for its clear channel list via a clear channel request message. In response to receiving a clear channel request message, a respective remote site transceiver 12 transmits back to the master the clear channel list it obtained during the preamble portion of the master site transceiver's message. The master site 10 continues to sequentially interrogate each of the remote transceivers, via subsequent clear channel list requests, until it has completed interrogation of every remote site transceiver.

**[085]** The master unit then logically combines all of the clear channel lists from all the interrogated remote transceivers to produce an 'aggregate' clear channel list. This aggregate clear channel list is stored in the

master site's transceiver and is broadcast to all of the remote site transceivers. When the aggregate clear channel list is received at a respective remote site it is stored in its transceiver.

**[086]** As noted earlier, both the master site transceiver 10 and the remote site transceivers 12 execute a synchronization function to ensure that the endpoint radios are synchronized to a common set of the clear channel list, channel hopping sequence, and preamble channel. A remote site transceiver 12 determines a lack of network synchronization by an absence of received network messages from the master. Primarily, a remote site transceiver will use the periodic aggregate clear channel list that it expects to receive from the master site as a network "heartbeat." A remote transceiver also employs a counter which indicates the amount of time elapsed from the last network message received. Then based upon the activity level of these two indicators, a remote site transceiver will determine whether it is properly synchronized with the network. If the remote site transceiver determines that it is not properly synchronized, it will fall into the normal initialization sequence for an uninitialized node. In the special case that a remote transceiver detects a primary user on the preamble channel, the remote site transceiver will send a message to indicate the conflict such that the master site transceiver 10 can indicate to the remaining remote transceivers 12 to fall back to the next preamble channel candidate based on a predetermined

algorithm. The remote site transceiver will then switch to the next preamble candidate as well. In the case where the master site transceiver detects a primary user on the preamble channel, it will initiate a switch to the next preamble channel candidate by sending a message indicator to all remote site transceivers. Those remote site transceivers not in conflict with the primary user will receive the message and switch. Those remote site transceivers that are in conflict will likely be unable to receive the message and will ultimately time out and declare a loss of network synchronization. At this point they will initiate a normal initialization sequence.

**[087]** Attention is now directed to Figure 2, wherein the overall architecture of a spectral reuse transceiver in accordance with the present invention is diagrammatically illustrated as comprising a transmitter section 100 and a receiver section 200. The input to the transmitter section 100 is a digital information packet supplied from an associated packet buffer 300 into which packets are coupled from an associated serial communication link and intended for transmission to a remote transceiver site. Conversely, the output of the receiver section 200 is a digitized data packet that has been transmitted from the remote transceiver site, the packet being coupled to the packet buffer 300. A communications controller 400 supervises the operation of the transceiver.

**[088]** In the digital signal processing path through transmitter section 100, a respective packet from the

packet buffer 300 is initially subjected to forward error correction (FEC) encoding by an FEC unit 101. As a non-limiting example FEC unit 101 may comprise a Viterbi encoder, which performs conventional convolutional encoding of an input packet. The output of FEC unit 101 is then coupled to a modulator 102. In accordance with a non-limiting, but preferred embodiment of the present invention, modulator 102 is implemented as a differential quaternary phase shift keyed (DQPSK) modulator.

**[089]** Advantages of using DQPSK include the fact that carrier phase lock is not required; non-coherent demodulation can be used. Also, in an OFDM system, phase uncertainty can be caused by either carrier phase offset or sample timing offset. A sample timing offset produces a phase offset on the carriers that is linearly increasing with the carrier number. DQPSK relieves some of the sample timing synchronization requirements. In addition, synchronization in most OFDM systems utilize pilot carriers during acquisition and tracking. It is preferable not to use such techniques in the spectral reuse transceiver of the present invention, which frequency hops only a limited number of data bearing carriers. This makes synchronization more complex. DQPSK removes these problems. A disadvantage of using DQPSK is that the non-coherent demodulation results in approximately 3 dB BER performance degradation. However, the transceiver is not power limited.

**[090]** Figure 3 is a dual graphical plot of the theoretical bit error rate (BER) curve 21 for DQPSK compared with the variation 22 for coherent QPSK. As shown therein, QPSK offers better performance. However, since the invention is not power limited, it is able to use increased gain with DQPSK to overcome the 3 dB reduction in performance in comparison with QPSK.

**[091]** Pursuant to a preferred implementation, DQPSK modulator 102 utilizes  $\pi/4$  DQPSK. This means that each input to modulator 102 will result in a phase shift in the constellation of  $\pi/4$ ,  $3\pi/4$ ,  $5\pi/4$ , or  $7\pi/4$  degrees. The constellation is shown in Figure 4. As shown therein, a transition from the point (0.7, 0.7), for example, may take any one of four paths, none of which passes through zero. The following equations are used to perform differential encoding.

$$I_{n+1} = I_n * iData_{n+1} - Q_n * qData_{n+1}$$

$$Q_{n+1} = Q_n * iData_{n+1} + I_n * qData_{n+1}$$

**[092]** These equations produce a transition from the current constellation point to the next point given the input bit pair ( $iData_{n+1}$ ,  $qData_{n+1}$ ).

**[093]** The I and Q outputs of the DQPSK modulator 102 are coupled to a symbol mapper 103, which maps the I and Q data from the modulator into the appropriate carrier(s) in accordance with a carrier hopping sequence 106. Mapping data onto the carriers requires knowledge of a 'clear channel' map and the hopping sequence. As pointed out briefly above, and as will be explained in detail below, a clear channel map is a compilation of all

current 'available' frequencies within the band of interest - here which ones of 480-6.25 KHz spaced channels that make up the band from 217 MHz to 220 MHz are determined to be 'unused' and are thus available for use by the spectral reuse transceiver. The communication control processor supplies a two dimensional multibyte array as a clear channel map to the mapper 103. Each bit in the mapping array indicates whether a respective channel is clear and thereby available for use, or is to be avoided, as will be described. The hopping sequence is determined by means of a pseudo-random sequence defined by a generator polynomial and a starting seed value.

**[094]** The control processor computes the sequence of available carriers that will be used to transmit FEC encoded data given the above information. This computation is performed off-line (i.e., not in real time as the message burst is transmitted). The maximal length packet is approximately 700 bytes. If, for example, 8 of 480 carriers are utilized to transmit the packet, and assuming that the FEC code rate is  $\frac{1}{2}$ , then 1 byte can be transmitted in one symbol. This results in a 700 symbol transmission, not including the preamble. The table may be shorter as the period of the PN sequence dictates. This assumes that each entry in the table is a carrier frequency index (stored as a multibit (e.g., 16 bit) value). During transmission of a packet, this look-up table (LUT) is used to select which carriers will be used. The same operation is performed in the



demodulator path of the transceiver, to derive which carriers have been received as will be described. For each pair of I/Q data retrieved from the DQPSK modulator 102, the symbol map look-up table is read to locate the appropriate carrier for this data. The value serves as an index into an array that is used to hold a 512 sample input vector for a downstream inverse Fourier transform (IFFT or  $\text{FFT}^{-1}$ ) unit 111 of modulated filter bank 110.

**[095]** The mapped carrier outputs of the symbol mapper 103 are coupled to a gain multiplier 104, which operates on a per transmission burst basis, and serves to minimize the power level to complete the radio link with a remote site, so as to avoid interfering with other links. For this purpose, the gain multiplier 104 is further coupled to receive a set of gain values stored in a gain value table 105, which is loaded by the control processor. Over time, statistics may be gathered on bit error rate (BER) to allow dynamic adjustment of the power to a remote site. Also, power is controlled within the burst, so that the preamble can be transmitted at maximum power (e.g., 2 W in one 6.25 KHz carrier). The data-bearing portion of the burst may be nominally transmitted at 0.25 W per subcarrier.

**[096]** The output of the gain multiplier 104 is coupled to a modulated filter bank 110, which is comprised of IFFT unit 111 and an associated polyphase filter 112. (It can be shown mathematically that a polyphase filter placed at the output of an inverse fast Fourier transform functions as a modulated filter bank.) As pointed out

above, in the exemplary embodiment described herein, the present invention operates within the 217-220 MHz frequency band, with adjacent carrier spacings of 6.25 KHz, yielding a total of 480 available carriers. However, among this substantial number of available frequencies only a relatively small fraction of the carriers (e.g., on the order of a dozen to several tens of channels) may actually be used.

**[097]** In order to accommodate the 480 carriers within the 217-220 MHz band of interest, IFFT unit 111 is preferably configured as a 512 point IFFT. Figure 5 depicts the multicarrier frequency spectrum produced by an IFFT absent the polyphase filter. As shown therein, for an adjacent pair of center frequencies 41 and 42, the sidelobe amplitudes are not constrained and occupy a substantial portion of the adjacent bandwidth, mandating the need for some form of spectral shaping (sidelobe containment). This is particularly true in the case of traditional OFDM, where the IFFT employs rectangular windowing in the time domain, resulting in a  $\sin(x)/x$  shaped frequency spectrum. The sidelobe performance of this function is relatively poor, with the first sidelobe being only 13 dB down from the main lobe. It may be noted that applications of OFDM are not concerned with the poor sidelobe performance of the IFFT output. This is because the traditional OFDM modems are designed to use the entire frequency band of interest (in the present example, 217 - 220 MHz). The sidelobes must be attenuated outside of the band of interest, and

this is done via a simple bandpass filter after the IFFT. In the present invention, there is no access to the entire band and each spectral reuse transceiver must avoid licensed users operating in the band of interest. Therefore, spectral shaping becomes critical.

**[098]** Although there are several ways to shape the spectrum of an OFDM modulator output, in accordance with the transceiver of the present invention, polyphase filtering is employed. It can be shown that a polyphase filter placed at the output of an IFFT behaves as a modulated filter bank. This means that while the polyphase filter is designed from a lowpass filter prototype, the polyphase filter becomes a bank of bandpass filters, which are frequency-shifted versions of the lowpass prototype. Figure 6 diagrammatically illustrates a mathematical model of a communication system employing such a modulated filter bank structure. As shown therein, an inverse discrete Fourier transform (DFT)-based transmitter 51 containing a lowpass filter operator  $H(z)$  is coupled to the transmit end of a communication channel 50, a receive end of which contains a DFT-based receiver 52 having a complementary lowpass filter operator  $H^{-1}(z)$ .

**[099]** In accordance with the invention, the lowpass prototype function  $H(z)$  in the mathematical model of Figure 6 is implemented as a six symbol wide interpolated or oversampled root-raised cosine filter, that is multiplied by a Tukey window (with  $\alpha = 5$ ) to further reduce sidelobes. As shown in Figure 7, the

lowpass impulse response of such a windowed filter mechanism settles rapidly to a relatively flat characteristic on either side of the main lobe, making it particularly suitable for use with an IFFT where only a subset of carriers are utilized and spectral leakage into adjacent carriers must be kept to a minimum. Such a windowed filter not only has superior sidelobe performance, but retains near zero inter-symbol interference (ISI).

**[0100]** The composite spectral response of such a modulated filter bank 110 for a pair of adjacent carrier frequencies is shown in Figure 8. As can be seen therein, as the carrier separation is increased to only two or three carrier increments, the sidelobe suppression is very substantial (on the order of 85-100 dB down). While the use of a modulated filter bank implemented by a cascaded arrangement of an inverse FFT and a polyphase for OFDM-based communications has been described in the literature (see for example, an article entitled: "Filtered Multitone Modulation for Very High-Speed Digital Subscriber Lines" by Giovanni Cherubini, et al., *IEEE Journal On Selected Areas In Communications*, Vol. 20, No. 5, June 2002), tailoring of the filter response characteristic as through the use of oversampling and application of a Tukey window, as well as the selective transmission of only a selected and varying subset of the available carriers, as carried out in accordance with the present invention, has not been described or suggested. (The Cherubini paper discusses

the motivation for Filtered Multitone Modulation as mitigating interference due to echo and near-end crosstalk.)

**[0101]** Namely, although the present invention employs the combination of an IFFT and a polyphase filter as an efficient way to partition the data onto a variable number of frequency hopped carriers, which is somewhat similar to OFDM, the spectral reuse transceiver of the invention does not utilize all of the possible carriers (512 for a 512 point IFFT). In fact, in the 217-220 MHz frequency band of the present example, there are only 480 available carriers. As pointed out above, of these, only a limited number subset are used.

**[0102]** Figure 9 diagrammatically illustrates the architecture of an interpolating or oversampled polyphase filter in accordance with the present invention. As shown therein, each filter section comprises a six stage delay line, respective  $z^{-1}$  stage outputs of which are coupled to coefficient multipliers, whose outputs are summed to realize an interpolated carrier value. In the present example, there is oversampling by 50%, namely, the interpolation factor is times 1.5, so that for a  $M=512$  point IFFT of the present example, the commutator 81 is rotated counterclockwise 768 times (i.e., for each symbol it steps through the 512 IFFT output samples 1.5 times), producing 768 samples. Thus, the sampling rate is increased by a factor of 1.5.

**[0103]** With the 512 complex values produced by the IFFT placed into six individual element delay line ( $z^{-1}$ ) stages, the contents of a given delay line form a vector of six complex elements, and there are 768 polyphase filter components computed from the prototype filter (shown in the Figure 9 as  $C_i$ ). Thus, each polyphase filter component contains six complex values. Each time the IFFT is executed, the commutator will cycle through all 768 component filters. When the commutator selects a given IFFT sample, the contents of a respective delay line are multiplied by the component filter (i.e., as a vector dot product) and then summed together to form one output sample of the polyphase filter.

**[0104]** The output of the polyphase filter is coupled to a times thirty-two interpolator 114. The use of this interpolator is based upon the parameters of the signal processing in the IFFT and the carrier spectrum of interest. In the present example, with a carrier spacing of 6.25 KHz, and an interpolation factor of 1.5, the effective IFFT rate is 6.25 KHz/1.5, or on the order of 4.1 KHz. The effective sample rate corresponds to the FFT rate, times the FFT size, times the interpolation factor, namely, the rate at which the complex data stream leaves the polyphase filter 112 is:

$$\text{Sample rate} = (6.25 \text{ KHz}/1.5) * 512 * 1.5 = 3.2 \text{ MHz.}$$

**[0105]** This data stream must be upsampled so that the data can be translated up in frequency digitally to an intermediate frequency of 70MHz. To this end, the 3.2 MHz data stream is digitally upsampled by a factor of

32, to realize an interpolated sample rate of 102.4 M samples/sec, at the output of interpolator 114. The output of interpolator 114 is then up-converted by digital mixer 115, which performs frequency translation to the intermediate frequency of 70MHz. The output of mixer 115 is then filtered by bandpass filter 117 and applied to digital-to-analog converter (DAC) 118. The output of DAC 118 is then translated via a frequency translator 119 to the desired 217-220 MHz transmission frequency band for transmission.

**[0106]** The receive signal path direction through the transceiver is generally complementary to the transmit path described above. In particular, a received signal in the 217-220 MHz transmission band is coupled to frequency translator 119, which down-converts the received signal to IF. The output of frequency translator 119 is coupled to analog-to-digital converter (ADC) 120, which produces a baseband digitized output that is applied to digital mixer 121. Digital mixer 121 down-converts the received signal further in accordance with the output of synthesizer 122 and applies the down-converted output to a decimator unit 123, the output of which is the same frequency as output by the polyphase filter 112, namely 3.2 MHz, as described above.

**[0107]** The received baseband data stream is then buffered in a timing correction buffer 124 and read out under the control of a symbol timing estimator 125 (to be described below with reference to Figure 12) and applied to a polyphase filter 126, the output of which is

coupled to an FFT unit 127. The output of the FFT unit 127 is coupled to three signal processing operators in the receive path. A first is a frequency offset estimator 128, the output of which is coupled to synthesizers 116 and 122, and which is operative to perform coarse and fine frequency estimation, as will be described. Coarse frequency estimation serves to reduce frequency offset to within a prescribed fraction (e.g., one-half) of adjacent carrier spacing; fine frequency estimation is performed on a burst-by-burst basis on the pure carrier portion of the preamble only. Fine frequency estimation is performed by looking at the change in phase of the carrier from symbol to symbol. The change in phase over four symbols is used as the fine frequency offset estimate.

**[0108]** A second signal processing operator, to which the output of the FFT unit 127 is coupled, is a symbol de-mapper 129. Symbol de-mapper uses the same mapping table 106 as the symbol mapper 103, so that the frequency data may be recovered. A third signal processing operator is a received signal strength indicator (RSSI) unit 130, which feeds the symbol timing estimator 125. The symbol timing estimator 125 is also coupled to the output of a mixer 133, which multiplies the received signal by the output of a preamble synthesizer 134, and is used to extract energy in the preamble that will enable the symbol timing estimator to perform its function.

**[0109]** Since the first part of the preamble of the received signal is pure carrier, which provides an



indication that the preamble is beginning, detecting energy in that frequency bin means that the receiver needs to be monitoring a burst, and the output of the RSSI unit 130 is used to enable the symbol timing estimator 125 to find the boundaries of a received symbol. In the parameters of the present example, 768 samples per symbol are being supplied to the receiver, but the symbol boundaries are unknown. The manner in which symbol timing is estimated by symbol timing estimator 125 (Figure 12) is described below. Once the symbol boundaries are known, the symbol timing estimator causes an appropriate shift of the data in buffer 124 as necessary to apply the received data one symbol at a time into polyphase filter 126.

**[0110]** The de-mapped frequency representative I and Q data values from the symbol de-mapper 129 are coupled to a demodulator 131 which, in the present example, is a DQPSK demodulator. DQPSK demodulation is a non-coherent scheme where the phase difference is derived from symbols  $(I_k, Q_k)$  and  $(I_{k-1}, Q_{k-1})$  for a given carrier, and requires storage of the previous (I,Q) data for each data bearing carrier. The algorithm to perform differential QPSK demodulation is given by the following equations:

$$I_{Data_{n+1}} = \text{sign}(Q_{n+1} * Q_n + I_{n+1} * I_n)$$

$$Q_{Data_{n+1}} = \text{sign}(Q_{n+1} * I_n + I_{n+1} * Q_n)$$

**[0111]** The output of DQPSK demodulator 131 is coupled to a Viterbi decoder 132, which recovers the original data and couples the data to the packet buffer.

**[0112]** As described briefly above, the symbol timing estimator 125 is operative to cause an appropriate shift of the data in the elastic buffer 124 as necessary to apply the received data one symbol at a time into polyphase filter 126. As noted previously, a symbol corresponds to 512 complex samples that result from one execution of the IFFT 111 in the transmitter path. If perfect symbol timing is achieved in the receive path, then the same 512 samples will be input to the FFT unit 127. A symbol timing error means that the 512 samples input to the FFT are some number of samples offset from the true transmitted symbol.

**[0113]** Estimation of symbol timing is critical, since most of the processing downstream of the FFT 127 requires that data input to the FFT 127 be symbol-aligned. If the symbol is not aligned, several performance degradations will occur. First, ISI will occur because the symbol presented to the FFT has samples from two different symbols. Also, a symbol timing offset will create a linearly increasing phase offset with respect to frequency in the post FFT data. Another issue for symbol timing recovery is adjacent channel interference (ACI). The symbol timing recovery algorithm must perform acceptably in the presence of ACI.

**[0114]** The BER curve plotted in Figure 10 indicates how sensitive the demodulator is to symbol timing errors. As seen in the plot, a symbol-timing offset of 16.7% of symbol duration will result in a performance degradation of approximately 2 dB. It is desired that symbol-timing

provide an estimate that will fall well within 16% of a symbol.

**[0115]** To aide in the recovery of symbol timing, the preamble of a respective transmission burst contains a plurality of N symbols of alternating +1s and -1s (where N is approximately 32), as shown in Figure 11. To convert the preamble channel down to baseband, the symbol timing estimator employs the signal processing flow shown in Figure 12. This allows lowpass filtering to attenuate ACI. The resulting signal is the alternating +1/-1 symbols degraded by noise, phase/frequency offset, and whatever ACI transits through the lowpass filter.

**[0116]** As shown in Figure 12, the output of mixer 133 is coupled through a low pass filter 141, which serves to reject an interference that might be caused by nearby primary user that is close to the preamble channel, to a moving average (MA) correlator 142. The MA is a window that is one symbol long (namely, 768 samples of I and Q data). The output of the MA correlator 142 is shown in Figure 13. Each minimum in Figure 12 represents a time when the MA window is centered on a symbol transition point. (Note that in Figure 13, the minima occur every 96 samples. To make simulation times reasonable, a 64-carrier system was simulated.) The output of the moving average correlator 142 is monitored at 143 over one symbol to find the minimum. The sample index of the minimum is stored at 144 and the average of N operations is provided as a symbol timing estimate at 145. As

pointed out above, this result is used to adjust or shift the data in the buffer 124 as necessary to provide symbol alignment.

**[0117]** Because, in the presence of noise, the sample index for the minima will vary, multiple symbols are averaged together to obtain the final estimate. Simulation results show that the symbol timing estimator performs well in noise. It should be noted that the transmit power of the preamble is 10 dB higher than the data bearing portion of the burst. During the preamble, only one carrier is being used to transmit data. Therefore, a full 2 W of output power can be transmitted on the preamble carrier. Figure 14 plots symbol timing estimate in the absence of adjacent channel interference and indicates an  $E_b/N_0$  of 15 dB during the preamble. This means that the  $E_b/N_0$  is 5 dB during the data bearing portion of the burst. In practice, the FCC limits the amount of power per channel to 2 W resulting in an  $E_b/N_0$  of 14 dB during the preamble.

**[0118]** Figure 15 plots the performance of the symbol timing in the presence of adjacent channel interference. From a comparison of the data plotted in Figure 15 with that of Figure 14 it can be seen that the symbol timing estimator performs very well in the presence of noise and implies that the low pass filter 141 of Figure 12 is critical to performance. The transfer function of a forty tap lowpass filter of Figure 12 is shown in Figure 16. It should again be noted that the simulation results are for a 64-carrier system so that simulation times are

reasonable. It will be readily appreciated that the filter characteristic of Figure 16 is more complex for a 512-carrier system, as the sampling rate is greater. Even though more coefficients are required, implementation of the symbol timing filter is readily achieved with a decimated sampling rate. This is acceptable since symbol timing estimate accuracy need only be within one-sixth of a symbol.

**[0119]** It may be noted that during a transmission burst, symbol timing will drift, which gives rise to the issue as to whether it is necessary to track symbol timing during the data-bearing portion of a burst. Analysis has revealed that given the maximal length of a burst and given the accuracy of the oscillators, there is no need to track. Once the symbol-timing estimate is computed and applied at the beginning of the burst, symbol timing does not need to be tracked for the remainder of the burst.

**[0120]** As pointed out above, the receive signal strength indicator (RSSI) unit 130, which is used to detect the pure carrier portion of the preamble, triggers the demodulator to begin processing an incoming burst (i.e. begin frequency offset estimation). The objective is to reliably detect the pure carrier portion of the burst preamble in as few symbols as possible. The detection process is directly linked to the RF front end AGC unit 135, since primary users will cause the AGC unit 135 to attenuate the desired SRT receive signal.

**[0121]** A block diagram of the RSSI unit 130 showing connectivity to the AGC circuit 135 is shown in Figure 17. As pointed out above, this unit serves to detect the pure carrier portion of the preamble. Known a priori is the carrier frequency (i.e., the identity of the frequency bin of the FFT corresponding to that carrier frequency). The input to the RSSI unit is derived from the IF front end, which is coupled to an AGC attenuator 201, the output of which is digitized by an ADC 202 and coupled through a digital down-converter (DDC) 203 to an FFT unit 204. The output of the DDC is coupled to an AGC control loop 205 for adjusting the AGC attenuator 201. It is also coupled to a normalizer path that includes a linear converter 206, the output of which is coupled through a 1/X divider feeding a multiplier/mixer 207 feeding a mixer 208. The output of the FFT unit 204 is also coupled to mixer 208.

**[0122]** Because the front end AGC impacts all carriers being received, including energy from unwanted interferers, it is necessary to perform an averaging operation over a plurality of symbols (e.g., four, in the present example). For this purpose, the output of mixer 208 is coupled to a four-symbol moving average (MA) filter, which serves to average over four symbols whatever energy is coming out of the preamble channel bin. The output of the moving average filter is then compared to a threshold in unit 211. Whenever the threshold is exceeded, the RSSI output is triggered.

**[0123]** Figure 18 graphically plots the RSSI normalized moving average. As shown therein during the presence of preamble carrier, there is a very noticeable rise in power level, which is detected as the beginning of a burst. As noted above, and as shown in Figure 17, the preamble carrier detector is coupled with the automatic gain control, which is monitoring the power in the entire frequency band. This power must not saturate the analog-to-digital converter 202 or else information will be lost. Therefore the AGC loop detects the input power and adjusts AGC attenuator 201, so that the ADC does not saturate. This AGC attenuation of the front end is effectively removed by the normalization circuitry described above, to which the AGC correction output is also fed.

**[0124]** The importance of normalization will be appreciated by reference to Figure 19, which shows the output of the moving average filter 210 of Figure 17, for the case that AGC normalization is not performed. From a comparison of Figures 18 and 19 it can be seen the AGC normalization serves to provide a very strong identification of average preamble power.

**[0125]** Figure 20 diagrammatically illustrates the configuration of the front end AGC unit. The front end 211 of the AGC unit takes the absolute value of the I and Q data and feeds the absolute value data to a moving average operator 222, which derives an average of one symbol (768 samples in the present example). This provides a metric of the power level. From this average

a desired value 223 is subtracted at 224 to provide an error value. This error value is then subjected to a fast attack, slow decay gain adjustment 225 at 226 and filtered at 227 to provide the desired AGC control output.

**[0126]** Figure 21 depicts the resulting time domain waveform, containing pure carrier at the beginning portion of alternating +1's and -1's, as described above. The large amplitude portion of the waveform constitutes an adjacent interferer. The issue is how well does the AGC operation attenuate the interferer. This is shown in the waveform in Figure 22, wherein the front end attenuator is driven down hard (fast attack) at the onset of the interference frequency, so as to prevent the ADC from saturating as described above. As soon as the adjacent channel interferer is gone the power ramps back up, as shown.

**[0127]** If the down-converter fails to accurately frequency translate the baseband data consisting of 512 carriers, there will be a performance degradation. Figure 23 shows what happens to BER performance when there is a frequency offset in the receiver. The figure shows the degradation caused by several frequency offsets defined as a percentage of the carrier spacing. To compensate for this problem, the invention conducts a coarse frequency estimation and a fine frequency estimation. The coarse frequency estimation serves to reduce frequency offset to within a prescribed fraction (e.g., one-half) of adjacent carrier spacing (which is 6.25 KHz



in the present example). Then, a fine frequency estimation can be made by operating the preamble data post FFT.

**[0128]** The coarse frequency estimation operation involves having the FFT monitor the beginning of the preamble during a modem training exercise when the spectral reuse transceiver is first deployed. While training, the unit does not receive data packets, but instead monitors the incoming bursts and perform coarse frequency offset estimation. The location of preamble is discovered via the network join process, but there may be a frequency offset. The energy in the preamble will appear nominally in the frequency bin in which it is expected. In this case there would be no frequency offset. However, part of the energy may show up in an adjacent channel. This indicates that the down-converter local oscillator needs to move in frequency, so that all the energy shows up in the expected bin of the receiver FFT. The demodulator looks for power in a small group of channels surrounding and including the known preamble channel. The power of the FFT in the receiver is utilized to locate the preamble. Average power levels are monitored in these channels (FFT bins). The objective is to find the preamble and adjust the digital down-conversion to move the preamble into the desired FFT bin. Success is declared when most of the signal power is confined to the desired FFT bin. It should be noted that both the remote receiver digital down-conversion and the remote transmitter digital up-conversion will be adjusted based

on this frequency offset. This will remove frequency offset in both the uplink and downlink.

**[0129]** Once coarse frequency acquisition is completed, bursts can be fully demodulated. At this point, fine frequency estimation is performed on a burst-by-burst basis on the pure carrier portion of the preamble only. Fine frequency estimation is performed by looking at the change in phase of the carrier from symbol to symbol. The change in phase over four symbols is used as the fine frequency offset estimate.

**[0130]** This is accomplished as shown in Figure 24 using a CORDIC operator 231 to calculate the phase angle  $\theta$  of the incoming I and Q pure carrier data. The CORDIC routine needs only to be invoked at the symbol rate (i.e., at 6.25 KHz/1.5 in the present example) and the resulting phase angle averaged over four symbols. The phase difference is computed at 232.

**[0131]** Figure 25 shows the accuracy of the fine frequency estimator for an  $E_b/N_0$  of 5 dB when averaging over four pure carrier symbols. Figure 25 shows a variance of 0.01 (0.01 x 6.25 KHz), which is shown by the diamond curve in Figure 23, referenced above. This is a one time estimate, that does not require tracking during the burst.

**[0132]** As pointed out above, the present invention makes use of a plurality of 'tickler' tones (e.g., three to five a priori known carriers) in the (217-220 MHz) band of interest that are transmitted simultaneously over a short duration (e.g., four symbols, as shown) for the

purpose of initiating an action in the receiving site's transceiver. The question is how to detect this energy. At the outset, it should be noted that which channels constitute the tickler carriers will have been defined by the master site and are given to the remote sites. As shown in Figure 26, the frequency information for those channels is extracted from the FFT. The absolute values of the I and Q subchannels are summed at 251 and then normalized at 252 using the AGC value as derived from the AGC loop described above. The normalized value is then subjected at 253 to a moving average over a prescribed symbol width (four symbols in the present example). The resulting average is then compared with a threshold at 254. If the threshold is exceeded an output is triggered indicating the detection of tickler tones.

**[0133]** The manner in a network comprised of a plurality of the spectral reuse transceiver's described above is configured, and communications are conducted between the master site and respective remote sites will now be described with reference to Figures 27-37. As pointed out earlier, the network of the present invention is comprised of a single master site-based spectral reuse transceiver and one or more (usually a plurality of) remote site-based spectral reuse transceivers, with all communications being conducted only between the master site and the remote sites; there are no communications between remote sites.

**[0134]** As shown in the state diagram of Figure 27, when a transceiver is powered up at state 271, the transceiver

is in idle mode, waiting for a command or interprocessor communication (IPC) from its associated control processor, as to whether it is to operate as a master transceiver or remote site transceiver. Where the transceiver is to operate as a master, it transitions to state 272, and the master transceiver proceeds to conduct network initiation and supervisory actions, as will be described. Where it is to operate as a remote site, the transceiver transitions to state 273, wherein it conducts a 'join the network' routine with the master site transceiver. As will be described below, and as indicated by the 'search for join preamble' loop 274, this involves a search for a prescribed 'beacon preamble' burst (to be described), which is periodically transmitted by the master for the purpose of enabling remote sites who wish to join the network to do so. Once the remote site has joined the network, it transitions to state 275, wherein it now has the ability to exchange data messages with the master, when polled by the master, as will be described.

**[0135]** As noted above, in accordance with the present invention, all actions, including the assembly of the communication network itself, are initiated by the master site transceiver. When the master site transceiver first comes up, it is the only member of the network. An initial task of the master is to determine whether there are any remote sites who wish to join the network, and then grant permission and enable such remote sites to become active network participants,

thereby assembling the network for its intended use (e.g., telemetry from a plurality of transducer sites). Once one or more remote site transceivers have joined the network, the master may transmit messages to those remote sites, and may grant permission to the remote sites to transmit messages back to the master site. To this end, the master site employs the four message formats shown in Figures 28-31.

**[0136]** More particularly, Figure 28 shows the contents of a 'beacon preamble' burst, that is periodically transmitted by the master for the purpose of stimulating a response from any remote site who wishes to join the network. To this end, the beacon preamble comprises a single carrier burst, a first portion 281 of which is pure carrier, which the master has determined after a scan of the spectrum of interest to be a clear channel. This clear channel carrier portion 281 is followed by a field 282 containing an alternating series of  $\pm 1$ 's, and terminated by a field 283, that contains a unique word specifically associated with a search for joining the network action. As will be described, in the course of scanning the (480) channels in (3 MHz) band of interest for the presence of activity, and detecting a beacon preamble, a remote site will proceed to transmit back to the master site a response burst containing only the carrier it has detected in the beacon preamble. The use of the carrier (which the master has previously determined to be a clear channel) in the beacon preamble

ensures that the response by the remote site will not interfere with another user of the network.

**[0137]** Figure 29 shows the contents of an initialization burst, which is transmitted by the master site to a remote site who is desirous of joining the network and has successfully responded to the master 'beacon preamble' shown in Figure 28, described above. Because the remote site has no knowledge of any clear channel other than the channel on which the master's beacon preamble was transmitted, it continues to listen on that channel for a follow-up initialization message from the master site. The follow-up or beacon message of Figure 29 is a single carrier message (the same clear channel which was detected by the remote site as the beacon of Figure 28) containing a preamble 291 of pure carrier, which is followed by a field 292 of alternating  $\pm 1$ 's, an a unique word field 293, which is different from the unique word field 283. This is followed by a message field 294, which contains prescribed information that enables the remote site to join the network, including the clear channel map, the PN sequence used to hop through the clear channel map, the seed for the PN sequence and the preamble channel number. As the remote transceiver is not locked to the master site transceiver, this last item ensures that the remote will properly identify the number of the channel on which it has responded to the master, and thereby enable the remote site to properly use the clear channel map for messaging.

**[0138]** Figure 30 shows the configuration of a standard data message burst used for the transmission of information between a master site and a remote site (other than initialization of the remote site, as described above, with reference to Figures 28 and 29). In particular, a data message burst contains a single channel preamble, an initial portion 301 of which is pure carrier, followed by an alternating series of  $\pm 1$ 's 302, and being terminated by a unique word field 303, that is different from the unique word fields of the message formats of Figures 28 and 29. The preamble, which may typically be on the order of several tens (e.g., 48) of symbols, is followed by a multicarrier data field 304 of N symbols in length.

**[0139]** Figure 31 shows the format of a 'tickler' burst. Respectively different sets of clear channels are used as tickler tones sets by the master site transceiver to initiate a prescribed response in a remote site transceiver, and by the remote site transceiver to initiate a response in the master site transceiver. In particular, as will be described, the master site transceiver may transmit a 'media open' tickler tone set to indicate that the network is available for the transmission of messages from a remote site transceiver to the master site; an 'access grant' tickler tone set granting access to the network to the first in time, access-requesting remote site transceiver; and a 'master access' tickler tone set to indicated to the network that the master site transceiver is about to broadcast a

message. A remote site transceiver may transmit an 'access request' tone set. This tone set is transmitted by a remote site having data to transmit to the master site transceiver, after the expiration of a random delay period following detection of the media open tickler tone set from the master site transceiver. Tickler tones may be comprised of sets of multiple frequencies (e.g., from three to five frequencies) extracted from the clear channel list and are transmitted simultaneously over a prescribed symbol span, e.g., on the order of four to five symbols.

**[0140]** Attention is now directed to Figure 32, which is a state diagram illustrating the manner in which a remote site transceiver may either join or reacquire the network. As shown therein, and as described above with reference to the master's 'beacon preamble' burst format shown in Figure 28, the master periodically transmits a 'beacon preamble', successful detection of which enables a remote site to join the network. To this end, in state 321, when a remote site is brought up or enabled by its control processor, it conducts a scan of the spectrum of interest for the presence of energy. As it scans the (480) channels it couples the received frequency energy information to its FFT in state 322. The outputs of the FFT are then examined in state 323 for the presence of potential candidates of a possible beacon preamble. Discriminating factors include the length of the burst (it being noted that a beacon preamble may be on the order of forty-eight symbols in length, so that a burst



on the order of forty-five to fifty symbols would make a good candidate), whether or not the detected energy was periodically received, and the power level of the received burst relative to other (data-conveying) carriers (a pure carrier containing a larger amount of energy than a data-bearing carrier).

**[0141]** Once the FFT has been examined and one or more carriers have been found, the received carriers are downconverted to baseband and examined in state 324 one at a time for the presence of the unique word associated with a beacon preamble, referenced above. If no candidate contains the prescribed unique word, the transceiver transitions back to state 322. If the proper unique word is found, however, the remote transceiver transitions to state 325, wherein it transmits back to the master site a response burst containing only the carrier it has detected in the beacon preamble, as noted above, the use of the carrier in the beacon preamble ensures that the response by the remote site will not interfere with another user of the network.

**[0142]** After transmitting the beacon carrier back to the master site transceiver, the remote site transceiver transitions to step 326, wherein it looks for the master to transmit the beacon burst of Figure 29. If the beacon burst of Figure 29 is received from the master within a prescribed time-out window, the transceiver transitions to state 327 wherein it joins the network. As described above, the message field of the beacon burst from the master site contains prescribed information that enables

the remote site to join the network, including the clear channel map, the PN sequence used to hop through the clear channel map, the seed for the PN sequence and the preamble channel number. On the other hand, if the beacon burst is not received from the master within the time-out window, the transceiver transitions to back to state 322.

**[0143]** As noted earlier, the preamble portion of a burst transmitted from the master site transceiver contains power from only a single carrier - making it the ideal time to conduct clear channel assessment (CCA), namely, the ideal time for each remote unit in the network to monitor the 3 MHz frequency band of interest (e.g., here the 217-220 MHz band) for the presence of energy (sourced from primary users). During the preamble, after received signal strength indication and frequency offset computations, most of the remaining preamble is used for symbol timing recovery. Therefore, during this initial time (until unique word processing begins), the primary goal of each unit's control processor is to monitor the 480-6.25 KHz channels in the 217-220 MHz band for the presence of interferers (primary users and others). Twenty-four to thirty-two symbols are available in each burst for clear channel assessment.

**[0144]** Clear channel assessment is performed at both the master site and at each of the remote sites and serves to maintain a current list of non-interfering (clear) channels at the master site and distributed to all remote sites in the network. By transmitting on only a

clear channel, a respective site is ensured that it will not interfere with any primary user of the spectrum of interest.

**[0145]** Figure 33 is a sequence diagram of the methodology through which the master site maintains and distributes this aggregate list of clear channels to all the remote sites in the network. When not transmitting a message to the master, each remote user is sequentially stepping through and monitoring its current list of clear channels (that it has previously obtained from the master unit), in accordance with a pseudo random hopping sequence known a priori by all the users of the network for a message that may be transmitted to it by the master site transceiver.

**[0146]** During the preamble period of any message being transmitted by the master at step 331, each remote node's transceiver scans all 480-6.25 KHz frequency bins within the 217-220 MHz spectrum for the presence of energy at step 332. Any bin containing energy above a prescribed threshold is masked as a non-clear channel, while the remaining ones of the 480 possible channels are marked as clear channels.

**[0147]** With each remote site transceiver having generated a clear channel list as a result of preamble scanning step 332, the master transceiver then sequentially interrogates each remote in the network for its clear channel list via a clear channel request message in step 333. In response to receiving a clear channel request message, a respective remote site transceiver transmits

back to the master channel at step 334 the clear channel list it obtained during the preamble portion of the master's message. The master site transceiver continues to sequentially interrogate each of the remote site transceivers, via subsequent clear channel list requests, until it has completed interrogation of the last remote site.

**[0148]** In step 335, the master site transceiver logically combines all of the clear channel lists from all the interrogated remote transceivers to produce an 'aggregate' clear channel list. This aggregate clear channel list is stored in the master transceiver and broadcast in step 336 to all of the remote transceivers. The aggregate clear channel list is broadcast to the remotes using the normal multi-carrier hopping transmission (not single-carrier). Only the initialization (beacon) message is transmitted on a single carrier. As the aggregate clear channel list is received at a respective remote site transceiver it is stored in memory.

**[0149]** As noted earlier, remote site transceivers are configured to conduct communications only with the master site. The master site employs an omnidirectional antenna, while the remote sites use directional antennas aimed at the master site. Communications are of two types: 1- the remote site has data to transmit and is awaiting permission from the master site to transmit that data to the master site; and 2- the master site has a message to transmit to a remote site. For the former

type of communication, the master makes the network or media 'open' to the first remote site transceiver to request access to the network in accordance with a prescribed 'random back-off' arbitration scheme. For the latter, the master broadcasts a prescribed 'master access' tickler burst.

**[0150]** The communication routine for the case that the remote site has data to transmit and is awaiting permission from the master site to transmit that data (to the master site transceiver) may be readily understood by reference to the state and sequence diagrams of Figures 34, 35 and 36. From an initial state 351 the remote site transitions to state 352 as it desires access to the network. In order to indicate that the network media is 'open' for message requests, the master site transceiver transmits a 'media open' tickler, shown at 361 in Figure 35 and 371 in Figure 36.

**[0151]** As shown at 371 in the contention and backoff diagram of Figure 36, each remote site transceiver with a pending message awaiting transmission will respond through a random slotted back off, before transmitting an access request - corresponding to the access request transmission 353 in the state diagram of Figure 34, the tickler 362 in Figure 35, and the access request 372 in the sequence diagram of Figure 36. Thereafter, the requesting remote transceiver waits for the master site to transmit an 'access grant' tickler, as shown at 363 in Figure 35 and 373 in Figure 36. Once a remote node has been granted access to the channel, as shown at

state 354 in the state diagram of Figure 34, the master node listens for a transmission (state 355 in the state diagram of Figure 34) from the remote node (data message 364 in Figure 35 and data message 373 in Figure 36) for a period of time shown as an acquisition of signal (AOS) timeout period 265 in Figure 35.

**[0152]** In the contention and backoff diagram of Figure 37, it can be seen that remote transceiver RTU2 will not attempt to send a data message, since it will not detect an access grant, as the access grant 373 from the master is transmitted at the same time that remote transceiver RTU2 is transmitting an access request. Remote transceiver RTU3 never attempts to send an access request, because it sees an access grant being transmitted by the master prior to RTU 3 initiating an access request, so that RTU3 knows that the access grant from the master site is intended for another remote transceiver.

**[0153]** Where the master site transceiver transmits a data message to a remote site, it transmits a prescribed master access tickler as shown at 366 in Figure 35. In response to this tickler, the remote site transceiver transitions from state 351 to state 356 in the state diagram of Figure 34. This is followed by the master site transceiver transmitting a message at 367 in Figure 35, which is received in state 356 in the state diagram of Figure 34.

**[0154]** The manner in which the master site communicates with remote sites may be readily understood by reference

to the state diagram of Figure 37. As shown therein the master site is initially in an idle state 350. In this state the master site transceiver is repetitively scanning the (480) frequency bins of the (3 MHz) bandwidth of interest, in order to update its clear channel list. It also uses a prescribed periodic time-out that controls the rate at which the master site transmits the beacon preamble of Figure 28 (step 381) for the purpose of allowing remote transceivers to join the network. It then transitions to step 382 wherein it listens for a response from any remote to the transmitted beacon preamble. If there is no response within a prescribed time-out window, the master transceiver transitions back to state 380. However, if a remote site transceiver response is received, the master site transceiver transitions to state 383, wherein it transmits the initialization beacon of Figure 29. The master site transceiver then transitions back to state 380.

**[0155]** If, when in idle state 380, the master site transceiver has a data message awaiting transmission, it transitions to state 384 wherein it transmits the master access tickler 366 of Figure 35. The master site transceiver then transitions to state 385 and transmits the data message, as shown at 367 in Figure 35. It then returns to the idle state 380.

**[0156]** In addition to enabling remote site transceivers to join the network and transmitting messages to remote site transceivers, the master site is also controllably

enabled to allow remote site transceivers to transmit data messages to it. For this purpose, as described above, the master site transceiver transmits a media open tickler as shown at 361 in Figure 35. In order to control when the media open tickler is transmitted, a media open time-out is employed for the case that the master site is in the idle state. When the media time-out tickler expires, the master site transceiver transitions to state 391, wherein it transmits the media open tickler. It then transitions to state 392 and looks for the return of an access request tickler from a remote site transceiver, as shown at 362 in Figure 35 and 372 in Figure 36.

**[0157]** If no access request tickler is received within a prescribed time-out, the master transceiver transitions back to state 380. On the other hand, where an access request tickler is received, the master transceiver transitions to state 393, wherein it transmits an access grant tickler, as shown at 363 in Figure 35 and 373 in Figure 36. It then transitions to state 394, wherein it receives the data message from the remote transceiver to whom transmission access of the network has been granted, as shown at 364 in Figure 35 and 374 in Figure 36. Once the data message has been received or a prescribed time-out expires, the master transceiver transitions back to idle state 380.

**[0158]** As will be appreciated from the foregoing description, the problem of how to allocate secondary usage of a licensed radio frequency band, so as not to



interfere with other users (licensed primary and secondary users), is successfully addressed in accordance with the present invention by means of a novel spectral reuse transceiver and an associated spectral activity-based link utilization control mechanism, that employ a selectively filtered form of orthogonal frequency division multiplexing to produce a sub-set of non-interfering radio channels.

**[0159]** While we have shown and described an embodiment in accordance with the present invention, it is to be understood that the same is not limited thereto but is susceptible to numerous changes and modifications as known to a person skilled in the art. We therefore do not wish to be limited to the details shown and described herein, but intend to cover all such changes and modifications as are obvious to one of ordinary skill in the art.